

1 CLAIMS:

2 What is claimed is:

3

4 1. A multi-service segmentation and reassembly (MS-SAR) integrated circuit,
5 comprising:

6 a first bus interface;

7 lookup circuitry;

8 segmentation circuitry;

9 reassembly circuitry;

10 a second bus interface; and

11 a data path extending from the first bus interface to the lookup circuitry, and from
12 the lookup circuitry to the segmentation circuitry, and from the segmentation circuitry
13 to the reassembly circuitry, and from the reassembly circuitry to the second bus
14 interface, wherein both cell-protocol traffic and packet-protocol traffic pass over the
15 data path from the first bus interface, through the lookup circuitry, through the
16 segmentation circuitry, through the reassembly circuitry and out of the integrated
17 circuit from the second bus interface, the lookup circuitry analyzing the cell-protocol
18 traffic and outputting information that causes the cell-protocol traffic to be processed
19 in a first way by the segmentation circuitry and the reassembly circuitry, the lookup
20 circuitry analyzing the packet-protocol traffic and outputting information that causes
21 the packet-protocol traffic to be processed in a second way by the segmentation
22 circuitry and the reassembly circuitry.

23

24 2. The integrated circuit of Claim 1, wherein the integrated circuit is operable in a
25 first ingress mode such that traffic is output from the integrated circuit to a cell-based
26 switch fabric via the second bus interface, and wherein the integrated circuit is
27 operable in a second ingress mode such that traffic is output from the integrated
28 circuit to a packet-based switch fabric via the second bus interface.

29

30 3. The integrated circuit of Claim 1, wherein the integrated circuit is operable in a
31 first egress mode such that traffic is received onto the integrated circuit from a cell-
32 based switch fabric via the first bus interface, and wherein the integrated circuit is

operable in a second egress mode such that traffic is received onto the integrated circuit from a packet-based switch fabric via the first bus interface.

4. The integrated circuit of Claim 1, wherein the integrated circuit is operable in an ingress mode such that traffic is output from the integrated circuit to a switch fabric via the second bus interface, and wherein the integrated circuit is operable in an egress mode such that traffic is received onto the integrated circuit from a switch fabric via the first bus interface.

5. The integrated circuit of Claim 1, wherein: 1) the integrated circuit is operable in a first ingress mode such that traffic is output from the integrated circuit to a cell-based switch fabric via the second bus interface, 2) the integrated circuit is operable in a second ingress mode such that traffic is output from the integrated circuit to a packet-based switch fabric via the second bus interface, 3) the integrated circuit is operable in a first egress mode such that traffic is received onto the integrated circuit from a cell-based switch fabric via the first bus interface, and 4) the integrated circuit is operable in a second egress mode such that traffic is received onto the integrated circuit from a packet-based switch fabric via the first bus interface.

6. The integrated circuit of Claim 1, wherein the cell-protocol traffic is ATM traffic, and wherein the packet-protocol traffic is MPLS traffic.

7. The integrated circuit of Claim 1, further comprising:
memory manager circuitry, wherein the data path extends from the segmentation circuitry to the reassembly circuitry via the memory manager circuitry.

8. The integrated circuit of Claim 1, wherein the cell-protocol traffic involves an ATM cell, and wherein the packet-protocol traffic involves a packet, the ATM cell being temporarily stored in one of a plurality of buffers of a memory, all of the buffers being of equal size, the packet being segmented into a plurality of chunks, and each of the chunks being temporarily stored into a corresponding one of the buffers.

9. An integrated circuit comprising:

a first bus interface;

means for generating a segmentation trailer;

means for checking a segmentation trailer;

a second bus interface; and

a data path extending from the first bus interface to the means for generating, and from the means for generating to the means for checking, and from the means for checking to the second bus interface, wherein both cell-protocol traffic and packet-protocol traffic pass over the data path from the first bus interface, through means for generating, through the means for checking, and out of the integrated circuit from the second bus interface,

10. The integrated circuit of Claim 9, wherein the integrated circuit is operable in an ingress mode and in an egress mode,

wherein in the ingress mode the integrated circuit is adapted for segmenting a packet into a plurality of segments, the means for generating a segmentation trailer generating a segmentation trailer and appending the segmentation trailer to one of the segments, the segments being output from the integrated circuit in the form of switch cells, and

wherein in the egress mode the integrated circuit is adapted for outputting packet information such that the packet information is transmitted as a packet onto a network, the means for checking receiving a plurality of segments, a last one of the plurality of segments including a segmentation trailer, the means for checking checking the segmentation trailer.

11. A switching device, comprising:

a first multi-service segmentation and reassembly (MS-SAR) integrated circuit;

a switch fabric; and

a second multi-service segmentation and reassembly (MS-SAR) integrated circuit, a flow of network information passing into the first MS-SAR, and then through the first MS-SAR, and then through the switch fabric, and then through the second MS-SAR, and then out of the second MS-SAR, wherein the flow passing into the first

MS-SAR is of a first traffic type, and wherein the flow passing out of the second MS-SAR is of a second traffic type, wherein the switching device can process the flow for all the four following pairs of first and second traffic types: 1) the first traffic type is ATM and the second traffic type is ATM, 2) the first traffic type is ATM and the second traffic type is packet, 3) the first traffic type is packet and the second traffic type is ATM, and 4) the first traffic type is packet and the second traffic type is packet, wherein the first and second MS-SAR integrated circuits are substantially identical integrated circuits.

12. The switching device of Claim 11, wherein when the first traffic type is ATM and the second traffic type is packet then the ATM traffic type involves AAL5 adaptation layer cells, and wherein when the first traffic type is packet and the second traffic type is ATM then the ATM traffic type involves AAL5 adaptation layer cells.

13. The switching device of Claim 11, wherein the switching device can also process a flow such that a single ATM cell is received onto the first MS-SAR and that ATM cell is output from the second MS-SAR encapsulated in a packet, there only being one ATM cell encapsulated in the packet.

14. The switching device of Claim 11, wherein the switching device can also process a flow such that a packet that encapsulates a single ATM cell is received onto the first MS-SAR, and wherein the ATM cell is de-encapsulated and output from the second MS-SAR as an ATM cell.

15. The switching device of Claim 11, wherein the switching device is an OSI layer three Internet Protocol (IP) router.

16. The switching device of Claim 11, wherein the switching device is an OSI layer two switch that does not perform Internet Protocol (IP) routing.

17. A multi-service segmentation and reassembly (MS-SAR) integrated circuit capable of processing a flow received from a switch fabric in accordance with a first

1 egress application type or in accordance with a second egress application type, an
2 indication of an application type being present in the flow as the flow is received onto
3 the MS-SAR, the MS-SAR locating the indication and if the indication indicates the
4 first egress application type then the MS-SAR processes the flow in accordance with
5 the first egress application type, but if the indication indicates the second egress
6 application type then the MS-SAR processes the flow in accordance with the second
7 egress application type.

8
9 18. The MS-SAR of Claim 17, wherein the flow is received on the MS-SAR in the
10 form of a switch cell, the switch cell including a switch header, the indication of an
11 application type being a plurality of bits in the switch header.

12
13 19. The MS-SAR of Claim 17, wherein the flow is received from one of a plurality of
14 input ports, each of the plurality of input ports having a port identification number
15 (port ID), the MS-SAR having, for each of the plurality of input ports, access to
16 locating information on where in a flow received on that input port the indication of
17 application type would be located, the MS-SAR using the port ID of a flow to access
18 the locating information, the MS-SAR using the locating information to locate in the
19 flow the indication of application type.

20
21 20. A method, comprising:

22 receiving on an integrated circuit from a switch fabric a first number of groups of
23 switch cells, each switch cell having a switch header and a data payload, each of the
24 groups being destined for one of a second number of output ports of the integrated
25 circuit, the first number being greater than the second number;

26 storing each of the data payloads of the first number of groups into a
27 corresponding one of a plurality of buffers;

28 retrieving from the plurality of buffers the data payloads of one of the groups of
29 switch cells destined for each of the output ports, and outputting from the integrated
30 circuit the data payloads retrieved such that a reassembled packet is transmitted
31 onto a fiber optic cable for each of the groups of switch cells retrieved, the
32 reassembled packet of a group comprising the data payloads of the switch cells of

1 the group, the data payloads of no more than one group for each output port being
2 output at any one time; and
3 maintaining a reassembly context for each group of data payloads being retrieved
4 from the plurality of buffers, each group of data payloads being retrieved being
5 destined for a different one of the output ports such that the integrated circuit
6 maintains no more than one reassembly context per output port.

7
8 21. The method of Claim 20, wherein the output ports are logical output ports
9 associated with a single physical output port.

10
11 22. The method of Claim 20, wherein the data payloads of one of the groups of
12 switch cells are AAL5 adaptation layer cells, one of the AAL5 adaptation layer cells
13 including a trailer, the trailer including a CRC, the integrated circuit not checking the
14 CRC in the trailer prior to storing the data payloads of the group into the plurality of
15 buffers.

16
17 23. A method, comprising:

18 receiving onto an integrated circuit from a fiber optic cable a first number of
19 packets, at most one of the packets being received at any one time on one of a
20 second number of input ports of the integrated circuit, the first number being greater
21 than the second number;

22 segmenting the packets on a per port basis such that at any one time at most one
23 packet is being segmented for each of the second number of input ports, the
24 segmenting of a packet resulting in a plurality of segments;

25 maintaining a number of segmentation contexts such that one segmentation
26 context is maintained per packet being segmented, and using the segmentation
27 context to generate a trailer that is appended to one of the segments of the packet
28 being segmented;

29 storing each of the segments of each of the first plurality of packets into a
30 corresponding one of a plurality of buffers, the trailer appended to a segment being
31 stored along with the segment into the plurality of buffers; and

1 retrieving from the plurality of buffers the segments of the first number of packets
2 and outputting each segment of the first number of packets from the integrated
3 circuit to a switch fabric in the form of a switch cell, wherein the number of
4 segmentation contexts maintained is less than the first number.

5
6 24. The method of Claim 23, wherein the input ports are logical input ports
7 associated with a single physical input port.

8
9 25. A system adapted for coupling to a switch fabric and a central processor, the
10 system comprising:

11 a first integrated circuit having a first bus interface, a second bus interface, and a
12 control interface, a data path extending from the first bus interface, through
13 segmentation circuitry on the first integrated circuit, through reassembly circuitry on
14 the first integrated circuit, and to the second bus interface;

15 a second integrated circuit that is substantially structurally identical to the first
16 integrated circuit, the second integrated circuit having a first bus interface, a second
17 bus interface, and a control interface, the second integrated circuit having a data
18 path extending from the first bus interface, through segmentation circuitry on the
19 second integrated circuit, through reassembly circuitry on the second integrated
20 circuit, and to the second bus interface; and

21 a control integrated circuit having a first control interface coupled to the control
22 interface of the first integrated circuit, having a second control interface coupled to
23 the control interface of the second integrated circuit, and having a third interface
24 adapted for coupling to the central processor, the control integrated circuit controlling
25 a first flow of network information out of the first integrated circuit, the first flow
26 passing over the data path of the first integrated circuit, the control integrated circuit
27 controlling a second flow of network information out of the second integrated circuit,
28 the second flow passing over the data path of the second integrated circuit, wherein
29 in operation either the first bus interface of the first integrated circuit and the first bus
30 interface of the second integrated circuit are coupled to the switch fabric or the
31 second bus interface of the first integrated circuit and the second bus interface of the
32 second integrated circuit are coupled to the switch fabric.

1
2 26. The system of Claim 25, wherein network information of the first flow is stored
3 by the first integrated circuit in a plurality of first buffers, all of the first buffers having
4 the same size, the control integrated circuit controlling the first flow of network
5 information out of the first integrated circuit by supplying an indication of a first buffer
6 to the first integrated circuit via the control interface of the first integrated circuit such
7 that the first integrated circuit retrieves the contents of the first buffer and outputs the
8 contents from the first integrated circuit, and

9 wherein network information of the second flow is stored by the second
10 integrated circuit in a plurality of second buffers, all of the second buffers having the
11 same size, the control integrated circuit controlling the second flow of network
12 information out of the second integrated circuit by supplying an indication of a
13 second buffer to the second integrated circuit via the control interface of the second
14 integrated circuit such that the second integrated circuit retrieves the contents of the
15 buffer and outputs the contents from the second integrated circuit.

16
17 27. The system of Claim 25, wherein the system can supply network information to
18 the switch fabric at a maximum system data throughput rate, and wherein the data
19 path through the first integrated circuit has a maximum data throughput rate, and
20 wherein the data path through the second integrated circuit has a maximum data
21 throughput rate, the maximum system data throughput rate being greater than the
22 maximum data throughput rate of the first integrated circuit and being greater than
23 the maximum data throughput rate of the second integrated circuit.

24
25 28. The system of Claim 25, wherein the system can receive network information
26 from the switch fabric at a maximum system data throughput rate, and wherein the
27 data path through the first integrated circuit has a maximum data throughput rate,
28 and wherein the data path through the second integrated circuit has a maximum
29 data throughput rate, the maximum system data throughput rate being greater than
30 the maximum data throughput rate of the first integrated circuit and being greater
31 than the maximum data throughput rate of the second integrated circuit.

32

1 29. An integrated circuit for processing network information, the integrated circuit
2 comprising a reassembly circuit, the integrated circuit being configurable into an
3 ingress mode and into an egress mode,

4 wherein in the ingress mode the reassembly circuit can use a flow identifier to
5 retrieve a switch header stored in a memory, the integrated circuit being adapted for
6 outputting a switch cell to a switch fabric in the ingress mode, the switch header
7 retrieved being included in the switch cell output to the switch fabric,

8 and wherein in the egress mode the reassembly circuit can use a flow identifier to
9 retrieve a network protocol header stored in a memory, the integrated circuit being
10 adapted for outputting either a cell or a packet to a network in the egress mode, the
11 network protocol header retrieved being included in the cell or packet output to the
12 network.

13
14 30. The integrated circuit of Claim 29, wherein the integrated circuit operates in the
15 ingress mode, wherein the flow identifier used to retrieve the switch header is
16 included in the switch header when the integrated circuit operates in the ingress
17 mode.

18
19 31. The integrated circuit of Claim 29, further comprising a mode control register,
20 wherein configuration information stored in the mode control register determines
21 whether the integrated circuit is configured in the ingress mode or is configured in
22 the egress mode.

23
24 32. The integrated circuit of Claim 29, wherein the integrated circuit operates in the
25 egress mode, the integrated circuit being capable of outputting both ATM cells and
26 MPLS packets to the network.

27
28 33. The integrated circuit of Claim 29, wherein the memory from which the switch
29 header is retrieved is a memory external to the integrated circuit, and wherein the
30 memory from which the network protocol header is retrieved is a memory external to
31 the integrated circuit.

32

1 34. The method of Claim 20, wherein the integrated circuit includes a reassembly
2 block, and wherein each reassembly context includes a running byte count value,
3 the running byte count value of the reassembly context for an output port being
4 updated as each successive data payload of the switch cells of a group destined for
5 the output port passes into the reassembly block, the reassembly block using the
6 running byte count value to remove padding from a last of the data payloads, the last
7 of the data payloads being marked by an EOP bit.

8
9 35. The method of Claim 20, wherein each reassembly context includes a running
10 byte count value and a partial CRC value.

11
12 36. A method, comprising:

13 outputting a plurality of chunks of information from an integrated circuit and
14 storing the plurality of chunks into a plurality of buffers in a payload memory external
15 to the integrated circuit;

16 retrieving the plurality of chunks from the payload memory and processing those
17 chunks sequentially through a reassembly block on the integrated circuit;

18 maintaining a running byte count value which is updated by the reassembly block
19 as each of the chunks passes into the reassembly block; and

20 using the running byte count value to reassemble the chunks such that the chunks
21 form a packet, wherein the integrated circuit is usable in an ingress mode and in an
22 egress mode, wherein in the ingress mode the packet is transferred to a switch
23 fabric, and wherein in the egress mode the packet is transmitted onto a network.

24
25 37. The method of Claim 36, wherein the packet is output from the integrated circuit
26 onto an SPI-4 bus, wherein the SPI-4 bus is coupled to the switch fabric if the
27 integrated circuit is operating in the ingress mode, and wherein the SPI-4 bus is
28 coupled to a framer if the integrated circuit is operating in the egress mode.

29
30 38. The method of Claim 36, wherein the chunks are processed through the
31 reassembly block sequentially such that a part of the packet is transferred out of the

reassemble block before one of the chunks has been retrieved from the payload memory.

39. An integrated circuit having one or more active output ports, comprising:
a memory manager that stores chunks of information into buffers in a payload memory, all the buffers being of equal size, some of the chunks including cell information, others of the chunks including packet information;
a reassembly engine that receives said chunks from the memory manager and that performs only one reassembly process at any given time for each active output port, the reassembly engine maintaining substantially no more than one reassembly context for each active output port, a reassembly context including a partial byte count and a partial CRC; and
a mode control register, wherein placing first configuration information in the mode control register causes the integrated circuit to operate in an ingress mode, and wherein placing second configuration information in the mode control register causes the integrated circuit to operate in an egress mode.

40. The integrated circuit of Claim 39, wherein the reassembly engine comprises:
a port calendar;
a data memory comprising a plurality of buffers;
an enqueue state machine for queuing said chunks into said data memory on a per output port basis; and
a dequeue state machine for dequeuing said chunks from said data memory in accordance with an output port identifier received from the port calendar.

41. An integrated circuit having one or more active output ports, comprising:
a memory manager that stores chunks of information into buffers in a payload memory, all the buffers being of equal size, some of the chunks including cell information, others of the chunks including packet information; and
reassembly means for receiving said chunks from the memory manager and for performing one reassembly process per active output port such that substantially no more than one reassembly context is maintained for each active output port.

1

2 42. The integrated circuit of Claim 41, wherein the memory manager stores said
3 chunks in the payload memory in per flow queues, and wherein the reassembly
4 means stores said chunks in per port queues.

5

6 43. The integrated circuit of Claim 41, further comprising:

7 a mode control register that stores configuration information, wherein the
8 integrated circuit is configured in an ingress mode if first configuration information is
9 stored in the mode control register, and wherein the integrated circuit is configured in
10 an egress mode if second configuration information is stored in the mode control
11 register.

12

13 44. The integrated circuit of Claim 41, wherein the reassembly means processes a
14 chunk in one of a plurality of ways as determined by type information, the type
15 information for a chunk being passed from the memory manager to the reassembly
16 means along with the chunk.

17